



10-16-00

IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE

PATENT APPLICATION

Ramesh Nagarajan
Muhammed A Qureshi
Steven A Surek

Express Mail® mailing label
number EJ049983137US
Date of Deposit OCTOBER 15, 2000

I hereby certify that this APPLICATION is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks Washington, D.C. 20231

K. Stone
(Printed name of person mailing paper or fee)
K. Stone
(Signature of person mailing paper or fee)

CASE 16-13-2

TITLE DATA COMMUNICATIONS NETWORKS WITH HIGH PERFORMANCE
AND RELIABILITYASSISTANT COMMISSIONER FOR PATENTS
WASHINGTON, D.C. 20231

SIR:

NEW APPLICATION UNDER 37 CFR § 1.53(b)

Enclosed are the following papers relating to the above-named application for patent:

Specification
14 Informal Sheets of drawing(s)
1 Assignment(s) with Cover Sheet
Declaration and Power of Attorney

| CLAIMS AS FILED | | | | |
|---|-----------|-----------|------------------|---------------|
| | NO. FILED | NO. EXTRA | RATE | CALCULATIONS |
| Total Claims | 38 - 20 = | 18 | x \$18 = | \$324 |
| Independent Claims | 11 - 3 = | 8 | x \$78 = | \$624 |
| Multiple Dependent Claims, if applicable | | | + \$260 = | \$0 |
| Basic Fee | | | | \$690 |
| | | | TOTAL FEE | \$1638 |

Please file the application and charge **Lucent Technologies Deposit Account No. 12-2325** the amount of \$1638, to cover the filing fee. Duplicate copies of this letter are enclosed. In the event of non-payment or improper payment of a required fee, the Commissioner is authorized to charge or to credit **Deposit Account No. 12-2325** as required to correct the error.

The Assistant Commissioner for Patents is hereby authorized to treat any concurrent or future reply, requiring a petition for extension of time under 37 CFR § 1.136 for its timely submission, as incorporating a petition for extension of time for the appropriate length of time if not submitted with the reply.



Please address all correspondence to **Docket Administrator (Room 3C-512), Lucent Technologies Inc., 600 Mountain Avenue, P.O. Box 636, Murray Hill, New Jersey 07974-0636**. However, telephone calls should be made to me at 732-949-1708.

Respectfully,



Joseph J Opalach
Reg. No. 36229
Attorney for Applicant(s)

Date: October 13, 2000

Lucent Technologies Inc.
600 Mountain Avenue (Room 3C-512)
P.O. Box 636
Murray Hill, New Jersey 07974-0636

DATA COMMUNICATIONS NETWORKS WITH HIGH PERFORMANCE AND RELIABILITY

FIELD OF THE INVENTION

This invention relates generally to communications and, more particularly, to
5 optical communications.

BACKGROUND OF THE INVENTION

Recently, it has been realized that the use of a signaling network is important for next generation intelligent optical networks for providing services like real time point-and-click provisioning of optical channels, optical layer protection and restoration, optical
10 layer network topology auto-discovery and optical layer bandwidth management. However, the performance and reliability of any signaling network (referred to herein as a control plane) is very important for successful deployment and usage of these services.

For a number of reasons, such as easier feature enhancement and wider access of features to customers, the Internet Protocol (IP) has been emerging as the technology of
15 choice to implement a control plane for Optical Transport Networks (OTNs). Unfortunately, the use of an IP-based control plane is not without its problems. For example, one form of routing used in an IP-based control plane is open shortest path first (OSPF) routing. Detection of a failure in OSPF relies on “hello” messages and the expiration interval of timers associated with these hello messages. Given this, time to
20 detect a failure is dictated by two factors. First, the frequency of hello messages, which must be computed while keeping in mind that resources, such as bandwidth and processing, required for them should not affect real traffic. Second, expiration values of timers, which must be chosen while realizing that congestion situations in the network should not be interpreted as failures. In practice, OSPF is slow to converge since the time
25 between hello messages and the expiration timer values, for a reasonable size network, are typically chosen on the order of seconds which makes the OSPF convergence time (failure detection + update of routing tables) on the order of several seconds (or sometimes maybe a few minutes). In case of a failure in such an IP-based control plane,

this slow convergence of OSPF routing degrades the setup time performance of new or existing connection requests in the system from on the order of tens of milliseconds to several seconds or minutes. Moreover, when the IP-based control plane failure is coupled with an OTN failure then restoration performance of the system gets severely impacted – 5 a situation that customers may perceive as intolerable.

SUMMARY OF THE INVENTION

As a result of the above, we have observed that failure detection in an IP-based control plane must not rely on the timer-based standard OSPF mechanism. Indeed, we propose an architecture that, in case of any failure in the IP-based control plane, can avoid 10 slow convergence of OSPF without compromising scalability. In accordance with the invention, a node of a control plane multi-feeds packets over diverse communications paths, and multi-selects packets received from diverse communications paths.

In an embodiment of the invention, an OTN (comprising a number of OTN nodes) uses an IP-based control plane (out-of-band signaling on a separate wavelength). Each 15 OTN node of the IP-based control plane performs dual-feeding and dual-selecting of signaling messages on diverse signaling paths. The IP-based control plane establishes a pair of physically disjoint signaling paths between every set of neighboring OTN nodes (pre-computed and pre-established physically disjoint primary and secondary message paths in the IP-based control plane).

20 BRIEF DESCRIPTION OF THE DRAWING

FIG. 1 shows an illustrative optical communications system embodying the principles of the invention;

FIG. 2 shows an optical communication system;

25 FIG. 3 shows an illustrative flow chart, embodying the principles of the invention, for use in the optical communications system of FIG. 1;

FIG. 4 shows another view of the optical communications system shown in FIG. 1;

FIGs. 5 and 6 show illustrative MLPS signaling between DCN nodes of the optical communications system shown in FIGs. 1 and 4;

FIGs. 7 and 8 show illustrative flow charts, embodying the principles of the invention, for use in a node of the optical communications system of FIG. 1;

FIGs. 9 and 11 show illustrative TCP connections in the context of dual-feeding and dual-selecting between nodes A and B of FIG. 4, in accordance with the principles of the invention;

FIGs. 10 and 12 show illustrative packets comprising sequence numbers;

FIG. 13 shows another illustrative network topology; and

FIG. 14 shows an illustrative high-level block diagram of a node in accordance with the principles of the invention.

10 DETAILED DESCRIPTION

Since an IP-based control plane is, in essence, another packet transport network for signaling messages, it is also referred to herein as a data communications network (DCN). As noted above, the DCN can have some failures of its own such as cable cuts. These may or may not be coupled with OTN failures depending on DCN and OTN topology. When the DCN has an independent failure, it may impact a new connection setup in the OTN as no signaling paths may be available in the DCN. When the DCN has a failure coupled with an OTN failure, it may not be possible to restore existing connections in the OTN even though alternate paths are available in the OTN. Clearly, this implies that the DCN has to be as reliable as the OTN itself and one needs fast restoration in the DCN itself in order to carry out the same for the OTN. Fast restoration in the DCN is quite challenging. However, a DCN can be constructed such that restoration time does not compromise OTN performance. In particular, and in accordance with the invention, a DCN node multi-feeds packets over diverse communications paths, and multi-selects packets received from diverse communications paths. For the purposes of the example described herein two diverse paths are used, i.e., a DCN node dual-feeds packets over diverse communications paths, and dual-selects packets received from diverse communications paths.

An illustrative optical communications system, in accordance with the principles of the invention is shown in FIG. 1. Other than the inventive concept, the elements

shown in FIG. 1 are well-known and will not be described in detail. For example, OTN 200 is an optical transport network comprising a number of optical cross-connect (OXC) nodes, e.g., OXC A, OXC B, OXC C, OXC D, OXC E and OXC F, having an illustrative OTN topology as shown. Also, although shown as a single block element, each node 5 (e.g., node A of DCN 100) includes stored-program-control processors, memory, and appropriate interface cards (not shown in FIG. 1). Except as noted below, it is assumed that OTN 200 conforms to a synchronous optical network (SONET). (It should be noted that other elements such as gateways to provide access to, e.g., OTN 200, and user endpoints, are left off to simplify the description.) In addition, the inventive concept uses 10 conventional programming techniques, which as such, will not be described herein.

As noted above, OTN 200 comprises OXC A, OXC B, OXC C, OXC D, OXC E and OXC F. The OTN 200 utilizes an IP-based control plane (out-of-band signaling on a separate wavelength) as represented by DCN 100. As such, DCN 100 comprises nodes 15 A, B, C, D, E and F. (In effect, this is a logical separation since each node – physically – performs both transport and signaling.) DCN 100 is a packet transport network for all the signaling messages necessary for connection signaling (e.g., setup and teardown), failure notification and OAMP (operations, administration, maintenance and provisioning) messaging in OTN 200. (Other than the inventive concept, signaling messages are known in the art and will not be described herein.) As will become apparent from the 20 description below, it is desirable that the DCN topology be independent of the OTN topology (however, it should be noted that common physical routes – while not required – are not excluded). DCN 100 utilizes any of a number of transport technologies such as, but not limited to, optical, SONET or Ethernet. This makes the DCN portable and applicable to any automatic switched transport network. FIG. 1 illustrates the inventive 25 concept for a sample connection setup in OTN 200 along path 201 (A-B-E-D). In addition, FIG. 1 shows the corresponding transport pathways in DCN 100. Note, that DCN 100 and OTN 200 do not share the same topology as DCN 100 is missing the OTN link B-E.

As noted above, the multi-feeding and multi-select operations require the DCN to 30 have physically diverse, pre-computed and pre-planned, paths – in this example, one

primary path and one secondary path. Physically diverse primary and secondary paths makes sure that in case of a failure in the DCN there is always a path readily available to exchange the control information (signaling). This thus enables the DCN to avoid slow OSPF convergence. However to compute physically diverse paths, the DCN needs 5 additional information about the mapping of IP (data) paths to the physical transport. For example, for the moment consider the DCN network illustrated in FIG. 2. In this figure, it can be observed that DCN node A has point-to-point links with DCN node B and DCN node C routed on the same physical link 51. This type of information is not available at the IP layer. Therefore if the DCN tries to compute a primary and a secondary path based 10 only on the IP layer information then it can end up computing paths that are not physically diverse. To make sure that the primary and secondary paths remain physically diverse, the DCN needs a mapping of its nodes to the underlying physical transport.

An illustrative flow chart for creating a DCN topology in accordance with the principles of the invention is shown in FIG. 3 for the optical communication system 15 illustrated in FIG. 1. Illustratively, two diverse paths are computed. In step 505, the OTN topology is identified in each DCN node. (For optical path computation purposes, OTN topology information is passed to each DCN node through a link state exchange protocol as known in the art (e.g., the Link Management Protocol (LMP)). In step 510, each DCN node identifies its neighboring DCN nodes according to the underlying OTN 20 topology determined in step 505. In step 515, each DCN node, for each of its identified neighbors, precomputes a primary and a secondary path to each of its identified neighbors (this information is stored in a routing table for each respective DCN node (not shown)).

Illustratively, multiprotocol label switching (MPLS) is used for the DCN network for explicitly routing control information along the precomputed paths. Unlike hop-by- 25 hop OSPF routing, MPLS cleanly separates the routing and the forwarding functions to provide explicit routing. Further, because of this separation, MPLS can easily isolate pre-computed MPLS paths (based on forwarding function) from any changes in the routing function because of a failure. (OSPF can also provide explicit routing capabilities through its explicit routing function but it requires that complete route must be carried in 30 each IP packet (waste of bandwidth). Also OSPF explicit routing is limited to path

lengths of nine hops that will put a constraint on number of hops in an optical path.)

As noted above, the DCN only provides a primary MPLS path and a secondary MPLS path between nodes that are neighbors according to the OTN topology. There are two reasons for this requirement. First, if primary and secondary MPLS paths are allowed between every pair of DCN nodes then on the order of n^2 bi-directional MPLS paths are required, where n is the number of nodes in the DCN network. This restricts the scalability of the DCN. Second, if the primary and secondary paths are allowed between DCN neighboring nodes, instead of DCN nodes that are neighbors according to the OTN topology, then the DCN network can become unnecessarily complex, which may also affect its performance. Given these two facts, if the MPLS paths are nailed between OTN neighbors then the control information needs to traverse only one label switched path (LSP) to go between two neighboring OTN nodes.

The OXC and DCN networks illustrated in FIG. 1 are redrawn in FIG. 4. In accordance with the inventive concept, and illustratively using MLPS, primary and secondary LSPs are required between neighboring OXC nodes: (A,B), (A,F), (E,F), (B,E), (B,C), (C,D), and (D,E). As mentioned above, the topologies of OTN 100 and DCN 200 are different, e.g., there is no DCN connection between nodes B and E. Now, suppose there is a request to setup an optical connection between OXC nodes A and D. Further suppose that A-B-E-D is the path computed for this connection request using the OTN topology (path 201 of FIG. 1). To setup this connection, signaling messages need to be sent between nodes (A,B), (B,E) and (E,D). Given LSPs between neighboring OXC nodes, messages will only have to traverse a single LSP between each pair of these nodes.

FIG. 5 illustrates an LSP between two neighboring OXC nodes, B and E, that are not neighbors in the DCN network (as shown in FIGs. 1 and 4). (Other than the inventive concept, the various protocols shown in the protocol stacks of FIG. 7 are well-known and will not be described herein. For example, consider the Transmission Control Protocol/User Datagram Protocol (TCP/UDP).) This LSP passes through intermediate DCN nodes C and D at the MPLS forwarding function level. (Note, that if MPLS paths are allowed between DCN neighbors then setup messages may need to travel many LSPs to go between two adjacent OXC nodes in the computed path. This is shown in FIG. 6,

which illustrates the sequence of LSPs signaling messages needed to travel between neighboring OXC nodes B and E in DCN network 100.)

In accordance with the inventive concept, each DCN performs a multi-select function (or mechanism or operation) and a multi-feeding (or replicating) function, which 5 occurs between every pair of adjacent OTN nodes. Referring back to FIG. 1, this is illustrated by node A and node B (multi-selection from, or multi-feeding to, other DCN nodes are not shown for simplicity). Illustratively, a duplicating function is performed. Node A performs dual-feeding of signaling packets to node B as represented by dotted arrows 101 and 102. Node B performs dual-selection of these signaling packets.

10 Two components required for any dual feed and dual select mechanism are: 1) the ability to dual feed at one end, and 2) the ability to select appropriately from the dual fed signal at the other end. A dual feeding method for use in, e.g., DCN node A, for transmitting a signaling packet is illustrated in FIG. 7. In step 205, DCN node A identifies two diverse signaling paths for transmission of the signaling packet to node B 15 (e.g., using information stored in the above-mentioned routing table (not shown)) – a primary path and a secondary path. Referring briefly back to FIG. 1, the primary path is the communications path (101) coupling DCN node A to DCN node B (i.e., they are adjacent in the topology shown in FIG. 1), and the secondary path is the communications path (102) through DCN nodes F, E, D and C. In step 210, node A assigns a sequence 20 number to the signaling packet. (Sequence numbers, by themselves, are known in the art. The range, or values, of sequence numbers is typically finite and repeats. For example, a 32 bit sequence number value can be used.) Every signaling packet receives a distinct sequence number from other signaling packets. In step 215, node A transmits the signaling packet to node B over the primary and secondary paths.

25 As the result of dual-feeding, at a destination node (such as DCN node B from the above dual-feeding example), each signaling packet is received at least twice. In the context of FIG. 1, DCN node B receives the signaling packet transmitted by DCN node A once from the primary path and once from the secondary path. For appropriately selecting the signaling packet exactly once, the destination node should be able to 30 identify the duplicate signaling packets and then select one. Any selection process should

take into account that duplicate signaling packets may not arrive at the same time (due to propagation delay and congestion) and also that these signaling packets may get lost (due to transmission errors and buffer overflows) – this is where the use of a sequence number in the signaling packet comes in. The use of a sequence number enables the destination

5 to identify duplicate signaling packets. (It should be noted that other forms of identification may also be used, if desired.) A dual-select method for use in, e.g., DCN node B, for receiving a signaling packet is illustrated in FIG. 8. In step 305, DCN node B receives a signaling packet comprising a sequence number value. The destination node, e.g., DCN node B, maintains a counter (or variable) (not shown), which indicates the sequence number value of the next signaling packet it is waiting for. Note, that node B keeps a separate counter for each DCN node to which it is a neighbor in the OTN topology. (It is assumed the counter value is initialized to a predefined starting value, e.g., 0.) In step 310, DCN node B checks the value of the received sequence number value against the value of the counter. (It should be noted that because of this sequence 10 number assignment scheme, for comparison purposes, the destination does not need to keep the first copy of any signaling packet until the second copy arrives.) If the value of the received sequence number is less than the value of the counter, DCN node B rejects the signaling packet on the basis that a copy of this signaling packet has already been received (from the other diversely routed path) in step 315 and returns to step 305 for receipt of the next signaling packet. On the other hand, when DCN node B receives a 15 signaling packet whose sequence number value equals the counter value, then DCN node B accepts that signaling packet and increases the counter value (e.g., by one) in step 320 and returns to step 305 for receipt of the next signaling packet. Finally, when the value of the received sequence number is greater than the counter value, DCN node B sets the counter value equal to the received sequence number value in step 325 and, in step 320, accepts that signaling packet and increases the counter value (e.g., by one), and returns to step 305 for receipt of the next signaling packet. This latter branch in the flow chart of FIG. 8 illustrates the loss of a signaling packet (whether significantly delayed or dropped). This puts responsibility on the higher protocol layers to recover from any out- 20 25 30 of-sequence/missing signaling packets. (Alternatively, other more complex procedures

may be used. For example, buffer the signaling packet and wait for the signaling packets with missing sequence number(s) to arrive from the other path and then accept all the signaling packets. The time to wait for a missing signaling packet must be limited to appropriately handle the scenario when both signaling packets (duplicate) with the same sequence number get lost. This choice tries to provide a reliable transport to the higher layers at the cost of more complicated implementation of the select function. It should be noted that the higher protocol layer is still responsible for the case when both signaling packets (duplicate) with same sequence number get lost.)

It should be noted that since the sequence numbers are allocated using a finite number of bits, they have a finite maximum value and thus will be repeated in module n fashion from 0 to $n - 1$. Given this, a mechanism is required to identify that a received packet with a sequence number less than the sequence number in the counter is either a duplicate packet or a new packet with a repeated sequence number (due to starting a new “round” of sequence numbers). It is suggested that the source pass additional information in the packet to help the destination identify the “round” of sequence numbers to which the packet belongs. This can be simply achieved by using a flag in the packet that changes when the source starts allocating sequence numbers again from the start after reaching the maximum allowable value. As such, the flow chart of FIG. 8 is suitably modified (not shown) so that the destination first checks the flag to determine if the sequence numbers have wrapped around. If the sequence numbers have not wrapped around, then the destination discards the packet with sequence number less value than the counter value (as shown in FIG. 8). However, if the sequence numbers have wrapped around, then the destination accepts the packet and updates the counter value.

Turning now to FIGs. 9 and 10, two illustrative approaches are shown for 25 realizing a DCN using dual-feeding and dual-selecting. FIG. 9 represents an illustrative implementation using two transmission control protocol (TCP) connections between each pair of DCN nodes that are neighbors according to the OTN topology. FIG. 9 illustrates this approach where nodes A and B are neighboring OTN nodes as shown in FIGs. 1 and 4. The dual feed and select function is implemented by dual feeding and then 30 appropriately selecting from the two TCP sessions which are diversely routed using the

two diversely routed LSPs. The packets from each TCP session are routed on the right LSP by using a FEC (Forwarding Equivalence Class) that comprises a destination IP address plus the TCP port number. Using retransmission and buffering, TCP provides a reliable transport that delivers packets in sequence. Signaling messages requiring reliable transport can access the DCN network through a defined application programming interface (API). (Note that in this case, the application running on top of TCP is responsible for assignment of sequence numbers on one end and selection from one of the dual-fed packets on the other end. An illustrative TCP packet comprising a sequence number (32 bits wide) is shown in FIG. 10.)

FIG. 11 illustrates another approach, which suggests using diversely routed LSPs with capability of dual feed and dual select. Note, that the MPLS does not have the capability to recover from loss of packets due to either transmission errors or buffer overflows. Therefore, the LSPs will see loss of packets. As noted above, implementing dual feed and select LSPs using the sequence number scheme puts responsibility of recovering from any missing/out-of-sequence packets on the above layers. To enable the higher level to recover from any packet losses, TCP sessions are used between neighboring OTN nodes. Specifically, one TCP session is established between each pair of DCN nodes, which are neighbors according to the OTN topology. Signaling packets using this TCP session are dual fed and dual selected, on to and from the diversely routed LSPs, at the ingress point and at the egress point, respectively, by the MPLS layer. Since MPLS does not provide dual feed and dual select functionality, it needs to be implemented at the MPLS layer. In the case when LSPs are established statically, conceptually, dual feed and dual select can be easily implemented by enhancing the FEC-to-label and label-to-FEC mapping functions. In the case when LSPs are dynamically established, enhancement in label distribution protocol (LDP) is also required. (Note, in this case the MPLS layer (below the IP layer) is responsible for assignment of sequence numbers on one end and selection from one of the dual-fed packets on the other end. An illustrative MPLS packet comprising a sequence number is shown in FIG. 12. As known in the art, an MPLS packet comprises an MPLS shim header 1. However, it should be noted that, and in accordance with the invention, an additional shim header 2 is used to

transport the sequence number following shim header 1. As such, the existence, and use, of an additional shim header to transport the sequence number is communicated between nodes during the setup portion of LDP.)

In light of the above, it should be noted that without the dual feed and dual select capability, if a failure occurs in the network, the DCN is required to detect for each LSP pair which of the LSP, primary or secondary, is affected by the failure. This means that the DCN needs to first find out where the fault has occurred (possibly in the OTN network) and then, if the DCN and OTN topologies are different, correlate the fault and the LSPs to determine per LSP pair which of the LSP has been affected. For example, consider the network shown in FIG. 13. Assume for the DCN portion of the network that A-B and B-C are primary LSPs with A-D-C-B and B-A-D-C as their diversely routed paths, respectively. Further suppose failure of the link between nodes A and B. It should be observed that this failure affects not only the primary LSP between nodes A and B but also the secondary LSP between nodes B and C. Note, that in absence of the above-described dual feed and dual select capability, the DCN needs to identify such distinctions which becomes complicated when the OTN and DCN topologies are not identical.

Turning briefly to FIG. 14, a high-level block diagram of a representative node 605 for use in accordance with the principles of the invention is shown. Node 605 is a stored-program-control based processor architecture and includes processor 650, memory 660 (for storing program instructions and data, e.g., for a routing table and for performing the above-described dual-feeding and dual-selecting illustrated in FIGs. 7 and 8, etc.) and communications interface(s) 665 for coupling to one or more communication paths as represented by path 666 (e.g., communication(s) interface 665 represents an optical dense wavelength division multiplexer (DWDM)).

As described above, the inventive concept provides a transport mechanism, for any IP-centric DCN network based, which can absorb a single failure in the DCN network with minimal, if any, impact on any signaling performance, connection setup time, and restoration time. In fact, one could state that failure detection or restoration is not required in such a DCN. However, in a more practical sense, these activities can be

carried out on a slower time-scale and do not impact real-time restoration of the OTN, which is of paramount interest. Another important characteristic of the inventive concept is that it does not limit, in any way, the scalability of the DCN network. Indeed, the inventive concept provides a reliable, high performance, DCN for next-generation
5 transport networks.

The foregoing merely illustrates the principles of the invention and it will thus be appreciated that those skilled in the art will be able to devise numerous alternative arrangements which, although not explicitly described herein, embody the principles of the invention and are within its spirit and scope. For example, although described in the
10 context of an IP controlled OXC-based optical transport network, the inventive concept is applicable to transport networks in general (utilizing an optical fabric and/or an electrical fabric) such as, but not limited to, PDH (Plesiochronous Digital Hierarchy); SDH (Synchronous Digital Hierarchy), Optical and other future transport network technologies. Also, although the inventive concept was illustrated in the context of a
15 primary path and a secondary path, the inventive concept is easily extended to any number of alternate paths. Finally, although illustrated in the context of an out-of-band signaling network, the inventive concept is applicable to an in-band signaling network as well.

WHAT IS CLAIMED:

1 1. A method for use in a node of a packet network, the method comprising the
2 steps of:

3 replicating a packet for transmission over at least two diverse communications
4 paths; and

5 multi-selecting a packet from at least two diverse communications paths.

1 2. The method of claim 1 wherein the packet is a signaling packet.

1 3. The method of claim 1 wherein the at least two diverse communications paths
2 used for replicating and multi-selecting are the same.

1 4. The method of claim 1 wherein the replicating step utilizes an Internet Protocol
2 (IP) for transmission over the at least two diverse communications paths.

1 5. The method of claim 1 wherein each replicated packet conveys an identical
2 packet identifier.

1 6. The method of claim 1 wherein each replicated packet conveys an identical
2 packet identifier in an additional shim header of a multiprotocol label switching (MPLS)
3 packet.

1 7. A method for use in a node of a packet network, the method comprising the
2 steps of:

3 identifying at least two diverse communications paths to an other node of the
4 packet network; and

5 replicating a packet destined for the other node of the packet network for
6 transmission over the at least two diverse communications paths.

1 8. The method of claim 7 wherein the packet is a signaling packet.

1 9. The method of claim 8 wherein the other node is a neighboring node in a

2 network topology of an underlying transport network.

1 10. The method of claim 8 wherein the replicating step includes the step of
2 inserting a sequence number into the packet for use by the other node in receiving the
3 replicated packets.

1 11. The method of claim 10 wherein the sequence number is inserted into an
2 additional shim header of a multiprotocol label switching (MPLS) packet.

1 12. The method of claim 8 wherein the packet is formatted in accordance with an
2 Internet Protocol (IP).

1 13. A method for use in a node of a packet network, the method comprising the
2 steps of:

3 receiving multiple copies of a packet; and

4 selecting one of the received multiple copies of the packet as a function of a
5 packet identifier in the received multiple copies of the packet.

1 14. The method of claim 13 wherein the packet is a signaling packet.

1 15. The method of claim 13 wherein the packet identifier is a sequence number

2 value.

1 16. The method of claim 15 wherein the selecting step includes the steps of:

2 if a value of the sequence number value of a received copy of the packet is less
3 than a counter value, discarding the received copy; and

4 if the value of the sequence number value of the received copy is equal to the
5 counter value, accepting the received copy.

1 17. The method of claim 15 wherein the selecting step includes the steps of:

2 if a value of the sequence number value of a received copy of the packet is less
3 than a counter value, discarding the received copy; and

4 if the value of the sequence number value of the received copy is not less than the

5 counter value, accepting the received copy.

1 18. A method for use in a node of a packet network, the method comprising the
2 steps of:

3 establishing a pair of physically disjoint paths to an other node that is a
4 neighboring node in a transport network; and

5 replicating a packet destined for the other node of the packet network for
6 transmission over the pair of physically disjoint paths.

1 19. The method of claim 18 wherein the replicating step includes the step of
2 inserting a sequence number into the packet for use by the other node in receiving the
3 replicated packets.

1 20. The method of claim 19 wherein the sequence number is inserted into an
2 additional shim header of a multiprotocol label switching (MPLS) packet.

1 21. The method of claim 18 wherein the packet is formatted in accordance with an
2 Internet Protocol (IP).

1 22. A method for use in a node of a packet network, the method comprising the
2 steps of:

3 establishing a pair of physically disjoint signaling paths to an other node that is a
4 neighboring node in an underlying transport network; and

5 replicating a signaling packet destined for the other node of the packet network
6 for transmission over the pair of physically disjoint signaling paths.

1 23. The method of claim 22 wherein the replicating step includes the step of
2 inserting a sequence number into the signaling packet for use by the other node in
3 receiving the replicated packets.

1 24. The method of claim 23 wherein the sequence number is inserted into an
2 additional shim header of a multiprotocol label switching (MPLS) packet.

1 25. The method of claim 22 wherein the signaling packet is formatted in
2 accordance with an Internet Protocol (IP).

1 26. A communications system comprising:

2 a number of nodes, each node storing at least two diverse communications paths
3 to a neighboring node, wherein the neighboring node is determined as a function of a
4 network topology of a transport network; and wherein each node, when transmitting a
5 packet to the neighboring node, replicates the packet for transmission over the at least
6 two diverse communications paths.

1 27. A communications system comprising a signaling network, the
2 communications system comprising:

3 a number of nodes, each node storing at least two diverse communications paths
4 to a neighboring node, wherein the neighboring node is determined as a function of a
5 network topology of an underlying transport network; and wherein each node, when
6 transmitting a signaling packet to the neighboring node, replicates the signaling packet
7 for transmission over the at least two diverse communications paths.

1 28. A communications system comprising:

2 a number of nodes, each node of the network performing dual-feeding and dual-
3 selecting of messages on diverse paths.

1 29. A communications system comprising a signaling network, the
2 communications system comprising:

3 a number of nodes, each node of the signaling network performing dual-feeding
4 and dual-selecting of signaling messages on diverse signaling paths.

1 30. A node for use in a network, the node comprising:

2 a memory for storing a routing table identifying at least two physically diverse
3 communications paths to a neighboring node; and

4 a processor for causing a packet to be replicated for transmission over the
5 identified at least two physically diverse communications paths to the neighboring node.

1 31. The node of claim 30 wherein the packet is a signaling packet.

1 32. The node of claim 30 wherein the packet is formatted in accordance with an
2 Internet Protocol (IP).

1 33. The node of claim 30 wherein each replicated packet conveys an identical
2 packet identifier.

1 34. The node of claim 30 wherein each replicated packet conveys an identical
2 packet identifier in an additional shim header of a multiprotocol label switching (MPLS)
3 packet.

1 35. A transmission frame representing a packet embodied in a transmission signal,
2 the transmission frame comprising:
3 a header; and
4 a packet identifier value for use in discriminating between multiple copies of the
5 packet.

1 36. The transmission frame of claim 35 wherein the packet identifier is conveyed
2 in an additional shim header of a multiprotocol label switching (MPLS) packet.


1 37. A transmission frame representing a signaling packet embodied in a
2 transmission signal, the transmission frame comprising:
3 a signaling message for use in establishing a connection in an associated transport
4 network; and
5 a signaling packet identifier value for use in discriminating between multiple
6 copies of the signaling packet.

1 38. The transmission frame of claim 37 wherein the signaling packet identifier is
2 conveyed in an additional shim header of a multiprotocol label switching (MPLS) packet.

ABSTRACT

An Optical Transport Network (OTN) (comprising a number of OTN nodes) uses an Internet Protocol (IP) based control plane (out-of-band signaling on a separate wavelength). Each OTN node of the IP-based control plane performs dual-feeding and 5 dual-selecting of signaling messages on diverse signaling paths. The IP-based control plane establishes a pair of physically disjoint signaling paths between every set of neighboring OTN nodes (pre-computed and pre-established physically disjoint primary and secondary message paths in the IP-based control plane).

FIG. 1 Nagarajan-Qureshi-Surek 16-13-2






FIG. 3 Nagarajan-Qureshi-Surek 16-13-2




FIG. 4 Nagarajan-Qureshi-Surek 16-13-2




FIG. 5 Nagarajan-Qureshi-Surek 16-13-2




FIG. 6 Nagarajan-Qureshi-Surek 16-13-2




FIG. 7 Nagarajan-Qureshi-Surek 16-13-2




FIG. 8 Nagarajan-Qureshi-Surek 16-13-2




FIG. 9 Nagarajan-Qureshi-Surek 16-13-2




FIG. 10 Nagarajan-Qureshi-Surek 16-13-2




FIG. 11 Nagarajan-Qureshi-Surek 16-13-2




FIG. 12




FIG. 13 Nagarajan-Qureshi-Surek 16-13-2





FIG. 14 Nagarajan-Qureshi-Surek 16-13-2



IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE

Declaration and Power of Attorney

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled **DATA COMMUNICATIONS NETWORKS WITH HIGH PERFORMANCE AND RELIABILITY** the specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by an amendment, if any, specifically referred to in this oath or declaration.

I acknowledge the duty to disclose all information known to me which is material to patentability as defined in Title 37, Code of Federal Regulations, 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

None

I hereby claim the benefit under Title 35, United States Code, 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

None

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint the following attorney(s) with full power of substitution and revocation, to prosecute said application, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected therewith:

| | |
|------------------------|------------------|
| Thomas J. Bean | (Reg. No. 44528) |
| Lester H. Birnbaum | (Reg. No. 25830) |
| Richard J. Botos | (Reg. No. 32016) |
| Jeffery J. Brosemer | (Reg. No. 36096) |
| Kenneth M. Brown | (Reg. No. 37590) |
| Donald P. Dinella | (Reg. No. 39961) |
| Martin I. Finston | (Reg. No. 31613) |
| William S. Francos | (Reg. No. 38456) |
| Barry H. Freedman | (Reg. No. 26166) |
| Julio A. Garceran | (Reg. No. 37138) |
| Jimmy Goo | (Reg. No. 36528) |
| Anthony Grillo | (Reg. No. 36535) |
| Stephen M. Gurey | (Reg. No. 27336) |
| John M. Harman | (Reg. No. 38173) |
| Matthew J. Hodulik | (Reg. No. 36164) |
| Michael B. Johannessen | (Reg. No. 35557) |
| Mark A. Kurisko | (Reg. No. 38944) |
| Irena Lager | (Reg. No. 39260) |
| John B. MacIntyre | (Reg. No. 41170) |
| Christopher N. Malvone | (Reg. No. 34866) |
| John F. McCabe | (Reg. No. 42854) |
| Scott W. McLellan | (Reg. No. 30776) |
| Martin G. Meder | (Reg. No. 34674) |
| John C. Moran | (Reg. No. 30782) |
| Michael A. Morra | (Reg. No. 28975) |
| Gregory J. Murgia | (Reg. No. 41209) |
| Claude R. Narcisse | (Reg. No. 38979) |
| Joseph J. Opalach | (Reg. No. 36229) |
| Neil R. Ormos | (Reg. No. 35309) |
| Jack R. Penrod | (Reg. No. 31864) |
| Gregory C. Ranieri | (Reg. No. 29695) |
| Scott J. Rittman | (Reg. No. 39010) |
| Ferdinand M. Romano | (Reg. No. 32752) |
| Eugene J. Rosenthal | (Reg. No. 36658) |
| Bruce S. Schneider | (Reg. No. 27949) |
| Ronald D. Slusky | (Reg. No. 26585) |
| David L. Smith | (Reg. No. 30592) |
| Ozer M. N. Teitelbaum | (Reg. No. 36698) |
| John P. Veschi | (Reg. No. 39058) |
| David Volejnicek | (Reg. No. 29355) |
| Charles L. Warren | (Reg. No. 27407) |
| Eli Weiss | (Reg. No. 17765) |

Please address all correspondence to the Docket Administrator (Rm. 3C-512), Lucent Technologies Inc., 600 Mountain Avenue, P. O. Box 636, Murray Hill, New Jersey 07974-0636. Telephone calls should be made to Joseph J Opalach by dialing 732-949-1708.

Full name of 1st joint inventor: Ramesh Nagarajan

Inventor's signature Ramesh Nagarajan Date 10/6/2000

Residence: Somerset, Somerset County, New Jersey

Citizenship: India

Post Office Address: 43 Oswestry Way
Somerset, New Jersey, 08873

Full name of 2nd inventor: Muhammad A Qureshi

Inventor's signature M. A. Qureshi Date 10/06/2000

Residence: Metuchen, Middlesex County, New Jersey

Citizenship: Pakistan

Post Office Address: 12 Carson Avenue
Metuchen, New Jersey, 08840

Full name of 3rd inventor: Steven A Surek

Inventor's signature S. A. Surek Date 10/13/2000

Residence: Leonardo, Monmouth County, New Jersey

Citizenship: United States of America

Post Office Address: 17 Koosman Drive
Leonardo, New Jersey, 07737